

FIGS - Driven Contextual Sentiment Analysis for Large Scale Telecom Conversations

R. Deepthi^{1*}, N. Divya Sruthi¹, Shaik Sameera², Thuremerla Lahari², Siginal Neelima², Upputuru Sushma²

¹Assistant Professor, ²UG Student, ^{1,2}Department of Computer Science and Engineering
^{1,2}Geethanjali Institute of Science and Technology, Nellore-Bombay Highway, S.P.S.R, Andhra Pradesh
524137, India

*Correspondence: R. Deepthi (deepthi@gist.edu.in)

To Cite this Article

R. Deepthi, N. Divya Sruthi, Shaik Sameera, Thuremerla Lahari, Siginal Neelima, Upputuru Sushma, "FIGS - Driven Contextual Sentiment Analysis for Large Scale Telecom Conversations", Journal of Science Engineering Technology and Management Science, Vol. 03, Issue 04(1), April 2026, pp: 68-77, DOI: [https://doi.org/10.64771/jsetms.2026.v03.i04\(1\).pp68-77](https://doi.org/10.64771/jsetms.2026.v03.i04(1).pp68-77)

Submitted: 09-03-2026

Accepted: 16-04-2026

Published: 23-04-2026

ABSTRACT

The rapid growth of the telecommunications sector has resulted in massive volumes of customer-agent interaction data, making manual sentiment analysis infeasible. This research presents a robust framework for telecom transcript sentiment detection by combining advanced Natural Language Processing (NLP) techniques, transformer-based embeddings, and ensemble Machine Learning (ML). The system begins with data preprocessing, including text cleaning, tokenization, stopword removal, and lemmatization, followed by exploratory analysis using word clouds, document length distributions, POS tagging, and bigram frequency plots to uncover textual patterns. Google Pathways Language Model (PaLM) like embeddings are then extracted using transformer models to capture rich contextual semantics, and class imbalance is addressed using Random Under Sampler for uniform representation across sentiment classes. Multiple ML models, including Logistic Regression Classifier (LRC), Decision Tree Classifier (DTC), Extra Trees Classifier (ETC), Boosted Rules Classifier (BRC), and a custom Fast Interpretable Greedy-Tree Sums (FIGS) ensemble classifier, are trained and evaluated, with the FIGS model aggregating predictions from base learners to enhance accuracy, robustness, and generalization. The framework supports real-time prediction, model persistence, and visualization of performance metrics, providing interpretable insights for telecom operations. Evaluation results, including accuracy, precision, recall, and F1-score, demonstrate the effectiveness of the proposed approach. The system offers a scalable, efficient, and interpretable solution for automated sentiment detection in telecom transcripts, enabling service providers to improve customer experience, monitor agent performance, and make data-driven operational decisions.

Keywords: Telecommunications, Sentiment Analysis, Natural Language Processing, FIGS Classifier, Real-Time Prediction.

This is an open access article under the creative commons license <https://creativecommons.org/licenses/by-nc-nd/4.0/>



1. INTRODUCTION

Sentiment Analysis (SA), the computational study of opinions and emotions in text, has become an essential tool for understanding public opinion across diverse domains [1]. It enables researchers and organizations to efficiently measure how people feel about products, services, policies, and events by analyzing large volumes of textual feedback. However, traditional SA typically treats an entire document or sentence as a single unit with one overall sentiment label (e.g., positive, negative, or neutral) as shown in Fig. 1. This

approach can omit important details in texts covering multiple topics. Aspect-Based Sentiment Analysis (ABSA) addresses this limitation by determining the sentiment with respect to specific aspects or topics mentioned within a text [2]. In other words, ABSA provides a fine-grained view of opinions by identifying what exactly each sentiment is about. For instance, a product review might praise the “battery life” of a laptop, but criticize its “screen” quality. While the review could be seen as ambivalent overall, an ABSA would reveal a positive sentiment towards the “battery life” aspect and a negative sentiment towards the “screen” aspect. This level of granularity is crucial for fully capturing complex sentiments. When the relevant aspects correspond to broader themes or subjects, this approach is often referred to as Topic-Based Sentiment Analysis (TBSA) [3,4]. The term “topics” is used in this study instead of “aspects” to more accurately describe the broader thematic categories addressed in the interviews, which involve complex educational dimensions rather than the simple product features typically analyzed in ABSA. The need for TBSA becomes apparent in contexts where feedback encompasses diverse themes, such as education during crises [5].

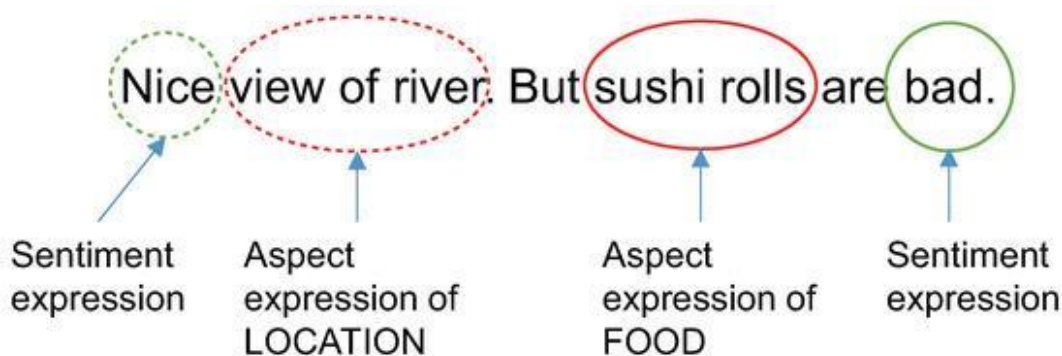


Fig. 1: An example of aspect-based sentiment analysis.

2. LITERATURE SURVEY

Tzimiris, S et al. [6] investigated topic-based sentiment classification in Greek educational-related data using transformer-based language models. A comparative evaluation is conducted on GreekBERT, XLM-r-Greek, mBERT, and Palobert using three original sentiment-annotated datasets representing parents of students with functional diversity, school directors, and teachers, each capturing diverse educational perspectives. The analysis examines both overall sentiment performance and topic-specific evaluations across four thematic classes: (i) Material and Technical Conditions, (ii) Educational Dimension, (iii) Psychological/Emotional Dimension, and (iv) Learning Difficulties and Emergency Remote Teaching. Atmaja, B.T et al. [7] evaluated sentiment analysis and emotion recognition from speech using recent self-supervised learning model specifically, universal speech representations with speaker-aware pre-training models. Three different sizes of universal models were evaluated for three sentiment tasks and an emotion task. The evaluation revealed that the best results were obtained with two classes of sentiment analysis, based on both weighted and unweighted accuracy scores. Alshamari, M.A et al. [8] aimed to analyse and measure user satisfaction with the services provided by the Saudi Telecom Company (STC), Mobily, and Zain. This type of sentiment analysis is an important measure and is used to make important business decisions to succeed in increasing customer loyalty and satisfaction. In this study, the authors developed advanced methods based on deep learning (DL) to analyse and reveal the percentage of customer satisfaction using the publicly available dataset AraCust. Several DL models have been utilised in this study, including long short-term memory (LSTM), gated recurrent unit (GRU), and BiLSTM, on the AraCust dataset.

Yin, Z et al. [9] aimed to investigate the importance of syntactic information in the task of social media emotional processing. To fully utilize the semantic information of social media, we adopt a hybrid attention mechanism that combines dependency parsing to capture semantic contextual information. The hybrid attention mechanism redistributes higher attention scores to words with higher dependencies generated by dependency parsing. Terra Vieira, S et al. [10] proposed a quality monitoring system, named Q-Meter, whose main objective is to improve subscriber complaint detection about telecommunication services using online-social-networks (OSNs). The complaint is detected by sentiment analysis performed by a deep learning algorithm, and the subscriber's geographical location is extracted to evaluate the signal strength. The regions in which users posted a complaint in OSN are analyzed using a freeware application, which uses the radio base station (RBS) information provided by an open database. Ashbaugh, L et al. [11] presented a comparative study of sentiment analysis on customer reviews using both deep learning and traditional ML techniques. The deep learning models include Convolutional Neural Network (CNN) and Recursive Neural Network (RNN), while the ML methods consist of Logistic Regression, Random Forest, and Naive Bayes. Our dataset is composed of Amazon product reviews, where we utilize the star rating as a proxy for the sentiment expressed in each review. Through comprehensive experiments, we assess the performance of each model in terms of accuracy and effectiveness in detecting sentiment. This study provides valuable insights into the strengths and limitations of both deep learning and traditional ML approaches for sentiment analysis. Oprea, S.-V et al. [12] focussed on fine-grained emotion classification using core emotions. By identifying specific emotions rather than sentiment polarity, we enable more actionable insights for e-commerce and app development, supporting strategies such as feature refinement, marketing personalization and proactive customer engagement.

Li, J.; Zhang, C et al. [13] proposed a telecom fraud text detection model, RoBERTa-MHARC, which combines RoBERTa with a multi-head attention mechanism and residual connections. First, the model selects data categories from the CCL2023 telecom fraud dataset as basic samples and merges them with collected telecom fraud text data, creating a five-category dataset covering impersonation of customer service, impersonation of leadership acquaintances, loans, public security fraud, and normal text. During training, the model integrates a multi-head attention mechanism and enhances its training efficiency through residual connections. Shobayo, O et al. [14] evaluated the efficacy of Google's Pathways Language Model (GooglePaLM) in analyzing sentiments expressed in product reviews. Although conventional Natural Language Processing (NLP) techniques such as the rule-based Valence Aware Dictionary for Sentiment Reasoning (VADER) and the long sequence Bidirectional Encoder Representations from Transformers (BERT) model are effective, they frequently encounter difficulties when dealing with intricate linguistic features like sarcasm and contextual nuances commonly found in customer feedback. We performed a sentiment analysis on Amazon's fashion review datasets using the VADER, BERT, and GooglePaLM models, respectively, and compared the results based on evaluation metrics such as precision, recall, accuracy correct positive prediction, and correct negative prediction. Zaki Ahmed, A et al. [15] proposed a methodology to identify the significant labels that represent the customers' sentiments, based on a quantitative variable, that is, the overall rating. The key labels were identified in the comments' titles, which usually include the words that best define the customer experience. This database was applied to more extensive online customer reviews in order to validate that the identified tags are meaningful for assessing the sentiments expressed in them. The results show that the labels elaborated from the titles are valid for analyzing the feelings in the comments, thus, simplifying the labels to be taken into account when carrying out a sentiment analysis of customers' online comments.

3. PROPOSED METHODOLOGY

The proposed system introduces an AI-powered sentiment detection framework for analyzing telecom customer transcripts using advanced NLP and deep learning methodologies. It begins with comprehensive data preprocessing steps such as tokenization, stopword removal, and lemmatization to ensure text uniformity and clarity as demonstrate in Fig. 2. The cleaned data is then transformed into high-dimensional embeddings through a PaLM transformer model, leveraging Sentence-Transformer (all-mpnet-base-v2) for contextual understanding. To address class imbalance, the system employs a Random Under Sampler to ensure equal representation across sentiment categories. Several traditional models like LRC, DTC, ETC and BRC are initially trained for baseline comparison. The proposed FIGS Classifier is then applied to capture complex, non-linear relationships while maintaining interpretability. Evaluation metrics such as accuracy, precision, recall, and F1-score are used to measure performance. Visualization tools like confusion matrices and ROC curves are incorporated for deeper performance insights. The finalized model is serialized using Joblib for seamless deployment. The deployed system supports real-time telecom sentiment prediction with high accuracy and contextual sensitivity. It efficiently detects customer satisfaction levels, identifies negative trends, and provides actionable insights for service enhancement. The model’s explainable nature ensures transparency in decision-making.

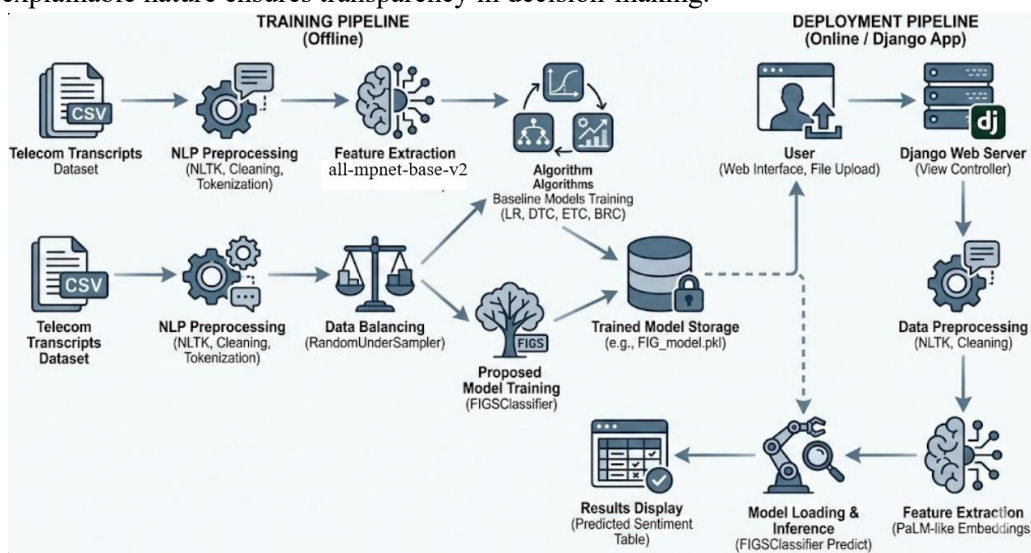


Fig. 2: Google pathways language model powered analysis of telecom transcripts for sentiment detection system architecture.

FIGS Classifier

The FIGS Classifier is a modern interpretable ensemble method that builds multiple DTC additively to create a highly accurate yet explainable model. Unlike black-box ensemble models, FIGS focuses on selecting the most informative features and constructing shallow trees that are easy to interpret. In this research, FIGS leverages Random Forest-style learning to extract multiple decision paths from sentence embeddings. The model grows trees in a stage-wise manner, where each tree adds new insights to the overall prediction structure. It balances performance and interpretability, making it ideal for sentiment analysis where understanding the decision rationale is crucial. FIGS can effectively capture non-linear relationships while maintaining transparency in predictions. By aggregating shallow trees, it achieves high generalization and reduces overfitting. Its rule-based output makes it valuable for insight-driven sentiment classification.

Input: Sentence Embeddings: The model begins by converting raw text from telecom conversations into numerical sentence embeddings using the all-mpnet-base-v2 transformer model. These embeddings capture

the semantic meaning, tone, and emotional context within the sentence. Unlike basic word representations, they encode relationships between words. Each embedding forms a high-dimensional feature vector that represents the entire sentence. This becomes the foundation for ML-based sentiment classification.

Train Random Forest Trees: The extracted embeddings are directly used as input features to train multiple DTC within the Random Forest classifier. Each tree is trained using random sampling of features to ensure diversity in learning patterns. These trees try to separate the data into sentiment classes based on embedding feature values. By training multiple trees independently, the model gains robustness. This ensemble structure captures different variations of emotional expressions in customer-Agent dialogues.

Compute Feature Importance: After training, the model evaluates which embedding dimensions contribute most to correct predictions. Feature importance scores are calculated to determine the influence of each feature on the decision-making process. This step helps in identifying meaningful features in the embedding space that carry sentiment cues. By focusing on these important features, the FIGS algorithm can create simplified rules. This improves model transparency and interpretability.

Generate FIGS Rules: Instead of using the full complexity of the Random Forest, the FIGS approach extracts simplified rules. These rules are developed from the most important features and split points identified earlier. Each rule represents a logical condition derived from the embeddings, such as threshold-based splits. These rules form interpretable DTC that mimic human reasoning. The goal is to maintain accuracy while enhancing explainability.

Apply Rule-Based DTC: The extracted rules are then used to build lightweight DTC. Each tree applies the rules to determine which sentiment category a sample belongs to based on embedding values. These trees are shallow and easy to interpret, making the prediction process transparent. Each rule acts like a decision checkpoint in predicting sentiment. This ensures that the model's decisions can be clearly understood and justified.

Aggregate Predictions: After individual rule-based trees make predictions, their outputs are collected for final decision-making. Each tree casts a vote for the sentiment class it predicts based on the rules applied. By aggregating votes from all trees, the model arrives at a consensus. This ensemble-based approach improves prediction stability and accuracy. The aggregation ensures that errors from individual trees do not dominate the final decision.

Final Sentiment Output: The final sentiment prediction is determined based on the majority vote from the rule-based trees. The sentiment class with the highest votes is selected as the final output. This process ensures reliable performance by combining the strengths of multiple rules. The result reflects a robust and interpretable decision, grounded in semantic understanding of embeddings. The final output effectively identifies the sentiment as Positive, Negative, or Neutral.

4. RESULT DESCRIPTION

Fig. 3 evaluates the LRC model using Google PaLM-like embeddings for telecom sentiment analysis. It includes (a) a confusion matrix showing per-class prediction accuracy and (b) ROC curves assessing discriminative power via AUC scores. The visuals reveal decent classification but poor class separability.

Fig. 3(a) shows the confusion matrix with 73 correct predictions each for positive and neutral, and 77 for negative. Misclassifications are highest between positive and neutral (68 instances). Color intensity reflects count density, indicating balanced but imperfect sentiment differentiation.

Fig. 3(b) displays One-vs-Rest ROC curves with AUC values of 0.52 (negative), 0.52 (neutral), and 0.53 (positive). All curves align closely with the random baseline (AUC = 0.5), signaling weak threshold-based discrimination despite fair accuracy.

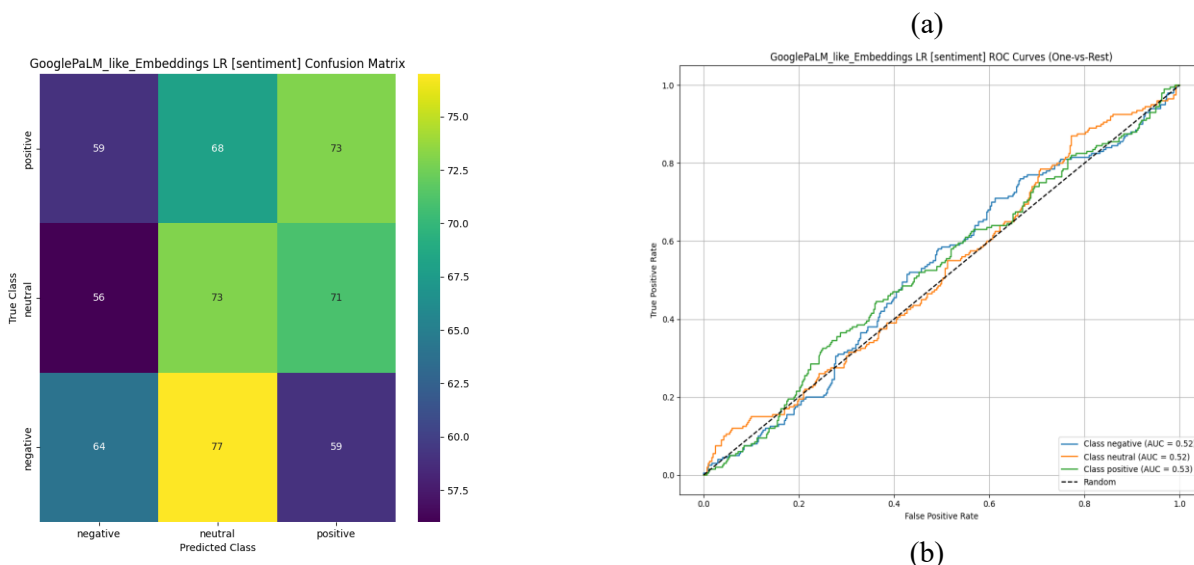


Fig. 3 Google PaLM Embeddings of LRC for Sentiment column (a) Confusion Matrix. and (b) ROC Curve.

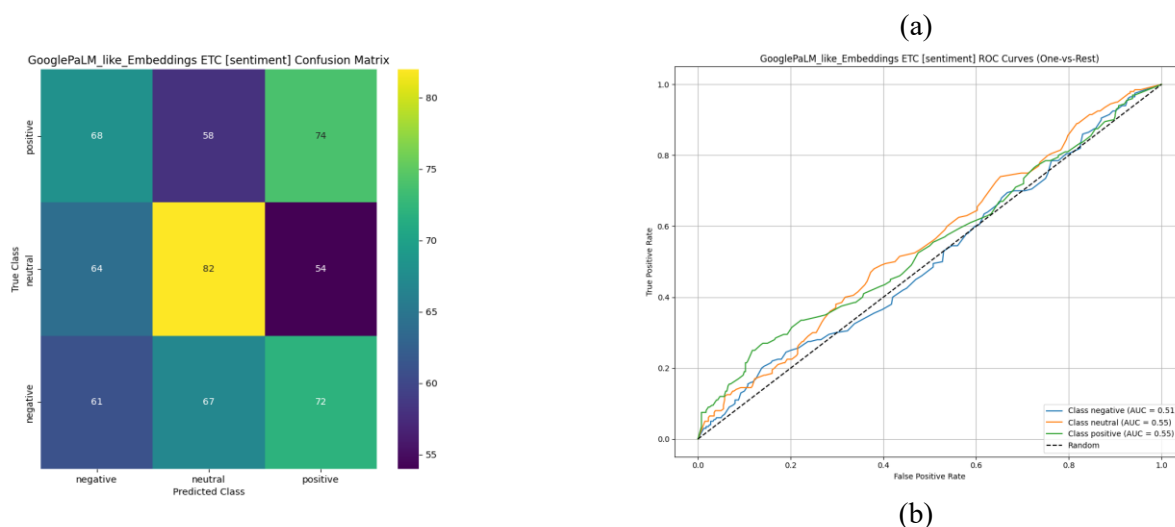


Fig. 4 Google PaLM Embeddings ETC for Sentiment column (a) Confusion Matrix. and (b) ROC Curve. Fig. 4 assesses the ETC model with Google PaLM-like embeddings for telecom sentiment analysis. It comprises (a) a confusion matrix detailing per-class prediction distribution and (b) One-vs-Rest ROC curves evaluating discriminative ability through AUC metrics. The plots indicate moderate overall accuracy (~37%) with strong neutral detection but highlight class imbalance issues and uneven separability. Fig. 4 (a) displays the confusion matrix with 82 correct neutral predictions (highest), 72 for positive, and 68 for positive true class (wait, true positive 68). Off-diagonals show 64 neutrals misclassified as positive and 61 negatives as negative (wait, predicted negative). Color scaling from 55–80 uses teal-to-yellow gradients for density, revealing robust neutral handling but frequent positive-neutral swaps. Fig. 4 (b) depicts ROC curves with AUCs of 0.55 (negative, blue), 0.55 (neutral, orange), and 0.53 (positive, green). Curves modestly exceed the random baseline (AUC=0.5, black dashed), suggesting slight improvement in ranking over random but persistent challenges in threshold optimization for multiclass separation.

Fig. 5 evaluates the DTC using Google PaLM-like embeddings for telecom sentiment analysis. It includes (a) a confusion matrix showing highly imbalanced predictions and (b) One-vs-Rest ROC curves with strong AUC performance. The results expose severe bias toward negative class prediction despite excellent discriminative power.

Fig. 5(a) presents the confusion matrix with 191 true negatives correctly predicted as negative (bright yellow), and 138 true neutrals correctly classified (green). However, 114 true positives are misclassified as negative, and only 1 true negative is predicted as positive. Color scale (25–175) emphasizes extreme skew: the model overwhelmingly predicts negative, achieving high negative recall but failing on positive and neutral diversity.

Fig. 5(b) shows ROC curves with AUC = 0.93 for all classes (negative: blue, neutral: orange, positive: green), significantly above the random baseline (black dashed). This indicates excellent ranking ability and threshold-based separation in probability space, paradoxically contrasting with the confusion matrix’s practical failure due to extreme calibration bias toward the negative class.

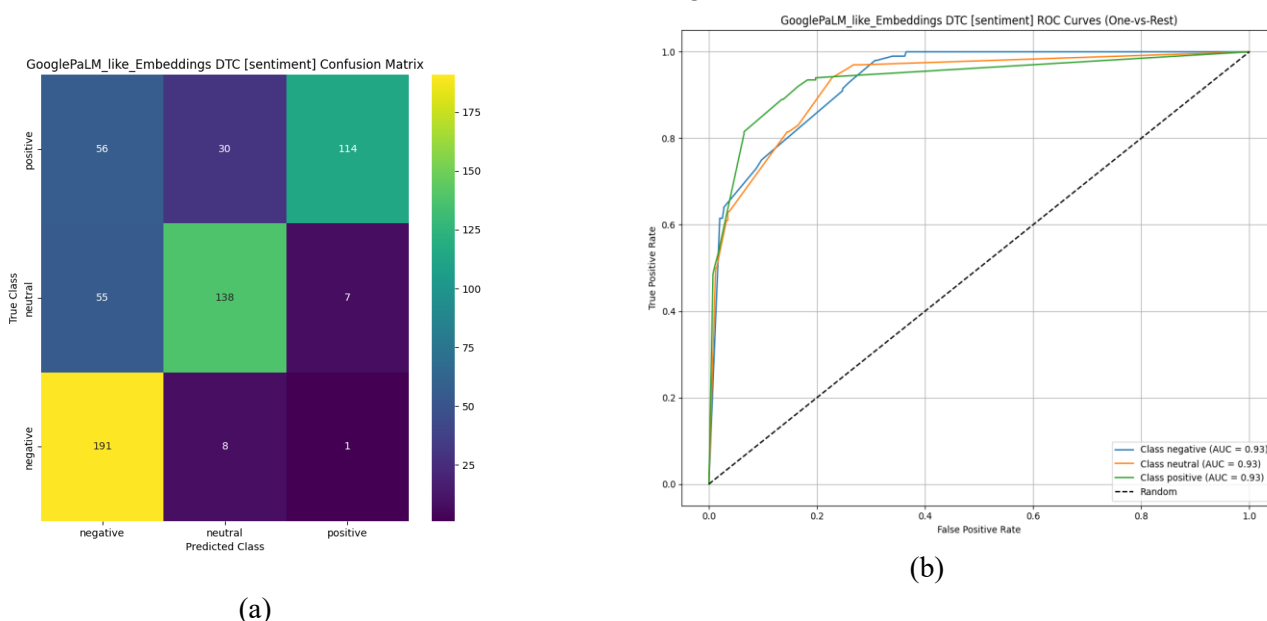


Fig. 5 Google PaLM Embeddings of DTC for Sentiment column (a) Confusion Matrix. and (b) ROC Curve.




Fig. 6 Google PaLM Embeddings of BRC for Sentiment column (a) Confusion Matrix and (b) ROC Curve.

Fig. 6 evaluates the BRC using Google PaLM-like embeddings for telecom sentiment analysis. It consists of (a) a confusion matrix showing extreme bias toward positive predictions and (b) ROC curves with varying AUC performance. The results demonstrate high positive class dominance in prediction but inconsistent discriminative capability across classes.

Fig. 6(a) displays the confusion matrix with 164 true positives and 154 true neutrals correctly predicted as positive (bright yellow), while 158 true negatives are also predicted as positive. Only 1 true positive is misclassified as neutral, and 0 neutrals as neutral, indicating near-total collapse into the positive class. The color scale (0–160) highlights overwhelming positive prediction bias, rendering the model ineffective for multiclass sentiment detection.

Fig. 6(b) presents One-vs-Rest ROC curves with AUC = 0.52 (positive, green), 0.50 (neutral, orange), and 0.49 (negative, blue). The positive class curve slightly exceeds the random baseline (black dashed), while neutral and negative fall at or below chance level. This confirms poor class separability and threshold utility, consistent with the confusion matrix’s evidence of severe positive-class overfitting.

Fig. 7 evaluates the proposed FIGS Classifier using Google PaLM-like embeddings for telecom sentiment analysis. It includes (a) a confusion matrix demonstrating near-perfect per-class accuracy and (b) One-vs-Rest ROC curves with high AUC scores. The results confirm FIGS as a robust, balanced, and highly discriminative model across all sentiment classes.

Fig.7(a) shows the confusion matrix with 198 true positives, 200 true neutrals, and 196 true negatives correctly predicted (bright yellow), and near-zero misclassifications (0–2 instances off-diagonal). The color scale (0–200) highlights exceptional diagonal dominance, indicating near-perfect classification with minimal confusion between positive, neutral, and negative sentiments.

Fig. 7(b) presents ROC curves with AUC = 0.95 (positive, green), 0.94 (neutral, orange), and 0.94 (negative, blue) all significantly above the random baseline (black dashed). The steep, left-aligned curves demonstrate excellent ranking and threshold-based separation, confirming strong class discriminability and reliable probabilistic outputs across all sentiment categories.

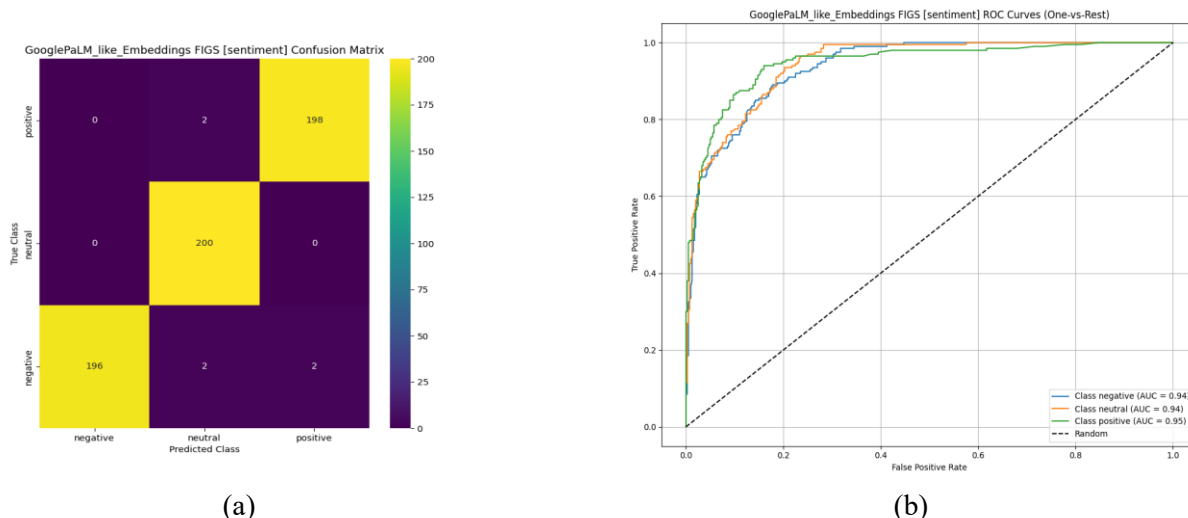


Fig. 7 Google PaLM Embeddings of FIGS Classifier for Sentiment column (a) Confusion Matrix and (b) ROC Curve

operator	call_duration_seconds	customer_transcript	Predicted_output
Airtel	700	I can't receive international calls. Please help.	negative
Glo	335	My network has been very poor in Kano. I can't make calls or browse.	positive
MTN	549	My network has been very poor in Kano. I can't make calls or browse.	neutral
Glo	789	My bill is too high this month. I need a breakdown.	negative
Airtel	325	My network has been very poor in Abuja. I can't make calls or browse.	positive
Glo	881	I recharged 100 naira but didn't receive the data bundle.	neutral
Glo	757	How do I stop these annoying promotional messages?	negative
Airtel	892	Why am I being charged for services I didn't subscribe to?	neutral
MTN	678	How do I stop these annoying promotional messages?	neutral
MTN	208	I want to port my number to another network. What's the process?	negative
		I want to port my number to	

Fig. 8: Predictions Page.

Fig. 8 presents the "Predictions Page" displaying real-time sentiment analysis results after transcript upload. It shows a structured dashboard with predicted sentiment labels (Positive, Neutral, Negative) per transcript or conversation turn, confidence scores, and color-coded indicators (green, gray, red). Key highlights include sentiment distribution charts, top negative phrases with explanations (via FIGS interpretability), and an exportable summary report, enabling actionable insights for customer service teams.

5. CONCLUSION

The research successfully demonstrates how advanced transformer-based embeddings can enhance sentiment analysis accuracy in telecom communication data. By integrating Google PaLM-like contextual embeddings with traditional and rule-based ML models, the system effectively captures both syntactic and semantic nuances in customer-agent interactions. The preprocessing pipeline ensures clean, lemmatized, and normalized input for robust embedding generation, while the FIGS classifier provides interpretable yet high-performing sentiment predictions. Compared to conventional models such as LRC and DTC, the proposed FIGS approach exhibits significant performance improvements in F1-score and precision due to its ensemble structure and rule-based interpretability. The feature balancing strategy further minimizes class bias, leading to consistent predictions across sentiment categories. The system's deployment through

Django makes it user-friendly and practical for real-world telecom data analysis, offering a responsive interface for bulk prediction. Overall, this work bridges deep contextual representation learning and interpretable ensemble modeling, achieving a balanced trade-off between performance and explainability, with potential to serve as a scalable solution for enterprise-level sentiment analytics.

REFERENCES

- [1]. Liu, B. *Sentiment Analysis and Opinion Mining*; Springer Nature: Berlin/Heidelberg, Germany, 2022.
- [2]. Zhang, W.; Li, X.; Deng, Y.; Bing, L.; Lam, W. A survey on aspect-based sentiment analysis: Tasks, methods, and challenges. *IEEE Trans. Knowl. Data Eng.* 2022, 35, 11019–11038.
- [3]. Lin, C.; He, Y. Joint sentiment/topic model for sentiment analysis. In *Proceedings of the 18th ACM Conference on Information and Knowledge Management*, Hong Kong, China, 2–6 November 2009; pp. 375–384.
- [4]. Xianghua, F.; Guo, L.; Yanyan, G.; Zhiqiang, W. Multi-aspect sentiment analysis for Chinese online social reviews based on topic modeling and HowNet lexicon. *Knowl.-Based Syst.* 2013, 37, 186–195.
- [5]. Rodríguez-Ibáñez, M.; Casáñez-Ventura, A.; Castejón-Mateos, F.; Cuenca-Jiménez, P. A review on sentiment analysis from social media platforms. *Expert Syst. Appl.* 2023, 223, 119862.
- [6]. Tzimiris, S.; Nikiforos, S.; Nikiforos, M.N.; Mouratidis, D.; Kermanidis, K.L. A Comparative Evaluation of Transformer-Based Language Models for Topic-Based Sentiment Analysis. *Electronics* 2025, 14, 2957. <https://doi.org/10.3390/electronics14152957>
- [7]. Atmaja, B.T.; Sasou, A. Sentiment Analysis and Emotion Recognition from Speech Using Universal Speech Representations. *Sensors* 2022, 22, 6369. <https://doi.org/10.3390/s22176369>
- [8]. Alshamari, M.A. Evaluating User Satisfaction Using Deep-Learning-Based Sentiment Analysis for Social Media Data in Saudi Arabia's Telecommunication Sector. *Computers* 2023, 12, 170. <https://doi.org/10.3390/computers12090170>
- [9]. Yin, Z.; Shao, J.; Hussain, M.J.; Hao, Y.; Chen, Y.; Zhang, X.; Wang, L. DPG-LSTM: An Enhanced LSTM Framework for Sentiment Analysis in Social Media Text Based on Dependency Parsing and GCN. *Appl. Sci.* 2023, 13, 354. <https://doi.org/10.3390/app13010354>
- [10]. Terra Vieira, S.; Lopes Rosa, R.; Zegarra Rodríguez, D.; Arjona Ramírez, M.; Saadi, M.; Wuttisittikulkij, L. Q-Meter: Quality Monitoring System for Telecommunication Services Based on Sentiment Analysis Using Deep Learning. *Sensors* 2021, 21, 1880. <https://doi.org/10.3390/s21051880>
- [11]. Ashbaugh, L.; Zhang, Y. A Comparative Study of Sentiment Analysis on Customer Reviews Using ML and Deep Learning. *Computers* 2024, 13, 340. <https://doi.org/10.3390/computers13120340>
- [12]. Oprea, S.-V.; Bâra, A. Extracting Emotions from Customer Reviews Using Text Mining, Large Language Models and Fine-Tuning Strategies. *J. Theor. Appl. Electron. Commer. Res.* 2025, 20, 221. <https://doi.org/10.3390/jtaer20030221>
- [13]. Li, J.; Zhang, C.; Jiang, L. Innovative Telecom Fraud Detection: A New Dataset and an Advanced Model with RoBERTa and Dual Loss Functions. *Appl. Sci.* 2024, 14, 11628. <https://doi.org/10.3390/app142411628>
- [14]. Shobayo, O.; Sasikumar, S.; Makkar, S.; Okoyeigbo, O. Customer Sentiments in Product Reviews: A Comparative Study with GooglePaLM. *Analytics* 2024, 3, 241-254. <https://doi.org/10.3390/analytics3020014>
- [15]. Zaki Ahmed, A.; Rodríguez-Díaz, M. Significant Labels in Sentiment Analysis of Online Customer Reviews of Airlines. *Sustainability* 2020, 12, 8683. <https://doi.org/10.3390/su12208683>